Unconditionally secure device-independent quantum key distribution with only two devices

Jonathan Barrett,[1,*] Roger Colbeck,[2,†] and Adrian Kent[3,4,‡]

[1] *Royal Holloway, University of London, Egham Hill, Egham TW20 0EX, U.K.*
[2] *Institute for Theoretical Physics, ETH Zurich, 8093 Zurich, Switzerland.*
[3] *Centre for Quantum Information and Foundations, DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge, CB3 0WA, U.K.*
[4] *Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON N2L 2Y5, Canada.*

(Dated: 11th October 2012)

Device-independent quantum key distribution is the task of using uncharacterized quantum devices to establish a shared key between two users. If a protocol is secure regardless of the device behaviour, it can be used to generate a shared key even if the supplier of the devices is malicious. To date, all device-independent quantum key distribution protocols that are known to be secure require separate isolated devices for each entangled pair, which is a significant practical limitation. We introduce a protocol that requires Alice and Bob to have only one device each. Although inefficient, our protocol is unconditionally secure against an adversarial supplier limited only by locally enforced signalling constraints.

Introduction

Key distribution is the task of establishing shared secret strings between two parties, and is sufficient for secure communication. Classical key distribution protocols base their security on assumptions about an eavesdropper's computational power. On the other hand, quantum key distribution protocols (e.g. [1, 2]) promise security against an arbitrarily powerful eavesdropper, and do so in the presence of realistic noise levels. However, in order for the security proofs to apply, the devices must operate according to certain specifications. Deviations from these can introduce security flaws, which can be difficult to identify (see e.g. [3] for practical illustrations of such attacks).

The difficulty associated with verifying the operation of quantum devices has led to much interest in device-independent quantum cryptography protocols. Ideally, such protocols guarantee security by tests on the outputs of the devices: no specification of their internal functionality is required. In a sense, the protocol verifies the devices' security *on-the-fly*.

Device-independent cryptography was first introduced by Mayers and Yao [4] (albeit under a different name) and the first quantum key distribution protocol to be proven device-independently secure was the Barrett-Hardy-Kent (BHK) protocol [5]. The BHK security proof applies not only against an arbitrarily powerful quantum eavesdropper (who also supplies the devices) but even against an eavesdropper and device-supplier who has discovered and makes use of any post-quantum physical theory, provided that, within the theory, the honest parties can enforce local signalling constraints. The applicability of the BHK protocol and proof to device-independent quantum cryptography was explicitly pointed out by later authors, who went on to develop some more efficient device-independent protocols with security proofs against restricted eavesdroppers [6–8] as well as other protocols shown to be unconditionally secure [9–13].

From a theoretical perspective, the BHK protocol provided an existence theorem for a task that had not been known to be possible. Practically, however, it has drawbacks. One is that, as formulated, it generates only a single bit of secure key. Although it can be modified using an idea from [14] to produce an arbitrarily long key, even with this modification, the protocol is inefficient and unable to tolerate reasonable levels of noise.

A serious practical problem with all the protocols with proven unconditional device-independent security [5, 9, 12, 13] is that they require that each (purportedly) entangled pair used in the protocol is isolated from the others. The protocols thus require a separate and isolated pair of devices for each entangled pair to ensure full device-independent security. This evidently makes such protocols costly to implement in practice.

We introduce here a protocol that evades this limitation, requiring only a single device for each user. Our protocol is a refinement of the BHK protocol, necessary in order to allow security when used with only two devices. As we have discussed elsewhere [15], the composability of device independent protocols is problematic if devices are reused in subsequent implementations. Here we show that if devices are not reused, then our protocol is secure according to a universally composable security definition, even against an adversary who supplies the devices and is restricted only by signalling constraints. As described, our protocol generates a single secure key bit. We also indicate how it can be modified using the idea in [14], to produce a key of arbitrary length. In addition, since it is composable, further key bits can be generated by running the protocol several times (although in this case, fresh devices are required for each run).

*jon.barrett@rhul.ac.uk
†colbeck@phys.ethz.ch
‡a.p.a.kent@damtp.cam.ac.uk

We see the value of our protocol as an existence theorem showing that device-independent quantum key distribution *is* in principle possible with only two devices. Whether this task can be achieved more efficiently and with reasonable noise tolerance remains (as far as we are aware) an open question.

We also show that some apparently natural extensions of existing protocols to two devices are insecure against eavesdroppers restricted only by signalling constraints, and in some cases also against quantum eavesdroppers. This may have impact in a recent line of work on the impossibility of privacy amplification against non-signalling eavesdroppers [16, 17].

Cryptographic scenario

We use a standard cryptographic scenario for key distribution. Here, two users (Alice and Bob), each have a secure laboratory in which to work, which they may partition into secure sub-laboratories. These allow Alice and Bob to prevent unauthorized communications between any devices they use. They are also each assumed to have (or be able to generate) their own supply of trusted random bits. To communicate between one another, Alice and Bob have access to an authenticated, but insecure, classical channel, and an insecure quantum channel. They may process classical information in a trusted way within their laboratories. However, any devices they use for quantum information processing are assumed to be supplied by an untrusted adversary (Eve). Eve may access (but not modify) any classical correspondence between Alice and Bob, and may access and modify quantum communication between them. She has complete knowledge of the protocol, but does not have access to the classical random data that Alice and Bob generate within their labs and use for the protocol (except for information she can deduce from what they make public).

Setup for the protocol

Alice and Bob each have a device, potentially supplied by Eve, that has an input port with $N \geq 2$ possible inputs and an output port with 2 possible outputs. Alice's inputs are denoted $A \in \{0, 2, \ldots, 2N-2\}$, and Bob's $B \in \{1, 3, \ldots, 2N-1\}$, and their respective outputs are denoted $X \in \{0,1\}$ and $Y \in \{0,1\}$. We define a set of allowed input pairs (A, B) by $\mathcal{G}_N := \{(0, 2N-1), (0,1), (2,1), (2,3), \ldots, (2N-2, 2N-1)\}$, with $|\mathcal{G}_N| = 2N$. For convenience, we introduce X' as a variable that is equal to $1 - X$ if $(A, B) = (0, 2N-1)$, and equal to X otherwise.

The devices are claimed by Eve to function by carrying out specified binary outcome measurements on the maximally entangled two qubit state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. Alice's input A is claimed to correspond to measuring the first qubit in the basis $\{\cos\frac{\theta}{2}|0\rangle + \sin\frac{\theta}{2}|1\rangle,\ \sin\frac{\theta}{2}|0\rangle - \cos\frac{\theta}{2}|1\rangle\}$, where $\theta = \frac{\pi A}{2N}$; similarly Bob's input B is claimed to correspond to measuring the second qubit in the basis defined by $\theta = \frac{\pi B}{2N}$. Alice and Bob do not need to test these precise claims, but instead perform various measurements and check their outcomes in such a way that the checks are unlikely to pass unless the produced bit is virtually as secure as a bit that would be generated were Eve's claims correct.

The protocol involves two security parameters: the integer $N \geq 2$ defined above, and a real number α in the range $0 < \alpha < 1$: to achieve reasonable security N needs to be large and α small. All classical communication between Alice and Bob is done via their authenticated classical channel.

Throughout the protocol, Alice and Bob keep their devices in isolated parts of their secure laboratories, ensuring that each device only learns its own inputs and cannot send any information outside the secure area. This ensures that the behaviour of the devices, which can be specified by a conditional probability distribution, satisfies certain non-signalling constraints. In particular, if the system Alice and Bob measure is correlated with a third system with input C and outcome Z, then the overall behaviour of the devices, $P_{XYZ|ABC}$, must be non-signalling, i.e. satisfy

$$\begin{aligned} P_{XY|ABC} &= P_{XY|AB} \\ P_{YZ|ABC} &= P_{YZ|BC} \\ P_{XZ|ABC} &= P_{XZ|AC}\,. \end{aligned} \tag{1}$$

These conditions ensure that if three parties possess devices with this behaviour, no subset of the parties can signal to any other subset by varying their choice of input.

Protocol R

1. Alice randomly chooses K, such that $K = 0$ with probability $1 - \alpha$ and $K = 1$ with probability α. She announces K to Bob.

2. On the i^{th} round, Alice picks a pair of values (A_i, B_i) at random from the set $\mathcal{G}_N$ specified above, and announces them both to Bob[1].

3. Alice inputs A_i into her device, and Bob B_i into his, and they record their outputs, the bits X_i and Y_i respectively. (Alice ensures that her device doesn't learn B_i.) If $(A_i, B_i) = (0, 2N-1)$, Alice sets $X'_i = 1 - X_i$, otherwise she sets $X'_i = X_i$.

[1] In fact, Alice need only announce B_i, but we have her announce both to make the analysis simpler.

4. If $K = 0$, Alice and Bob announce X'_i and Y_i. If $X'_i \neq Y_i$, they abort. Otherwise, they return to Step 1.

5. If $K = 1$, write $i = f$ (the final value of i). The bits X'_f and Y_f are taken to be the final shared secret key bit.

As presented above, this protocol requires Alice's and Bob's devices to contain sufficient pre-shared entanglement before the protocol starts. Taken literally, this requires an infinite supply of pre-shared $|\Phi^+\rangle$ states. More realistically, it requires a large number $M \gg \alpha^{-1}$ of pre-shared $|\Phi^+\rangle$ states, and that the parties accept a small probability of the protocol aborting because the supply is exhausted. These stringent technological requirements can be avoided by introducing an additional (untrusted) state-creation device, which could be incorporated into Alice's or Bob's measurement device, and which is supposed to generate $|\Phi^+\rangle$ states and send one qubit over the insecure quantum channel to the other party. The i^{th} state must be distributed before any information about the measurements (A_i, B_i) or the value of K is announced. This modification (call it Protocol R$^+$) gives Eve more cheating strategies but, as we show below, is still secure.

Security – main idea

The idea behind the security of this protocol is as follows. If the states and measurements are as Eve claims, then the quantity I_N defined by

$$I_N = I_N(P_{XY|AB}) := P(X = Y|A = 0, B = 2N-1) + \sum_{\substack{a,b \\ |a-b|=1}} P(X \neq Y|A = a, B = b)$$

satisfies

$$I_N = I_N^{\text{QM}} := 2N \sin^2 \frac{\pi}{4N} < \frac{\pi^2}{8N}\,. \qquad (2)$$

As N increases, these correlations give larger violations the chained Bell inequalities [18, 19], which in this formulation are $I_N \geq 1$.

The significance of this violation of the chained Bell inequalities for secrecy is that, in the limit of large N, the correlations that achieve the quantum bound (2) become monogamous and uniform [5, 14]. That is, for any non-signalling distribution $P_{XYZ|ABC}$ for which $I_N(P_{XY|AB})$ is small, and for any choice of input c, the outcome Z is virtually uncorrelated with X, and $P_{X|A=a}$ is virtually indistinguishable from uniform, for all a.[2] In other words, if Alice's and Bob's systems have a low I_N, then Eve (who we can take to hold the system with input C and output Z) must have almost no information about the outcomes they obtain. The protocol is designed so that (roughly speaking) if Eve supplies states for which there are many rounds in the protocol where I_N is high, the protocol is likely to abort, while if she supplies a state that has high I_N on only a few rounds, the round at which Alice and Bob finally (hope to) create the key bit is likely to have low I_N, and so the key bit is likely to be indeed both agreed by Alice and Bob and secure against Eve.

Our main result is that, if we choose $\alpha = N^{-\frac{3}{2}}$, and take N to be large, Protocol R is unconditionally secure, in the sense that the key bit it generates can be treated as though produced by a secure random key distribution oracle. Provided that the devices are not reused and are securely isolated, so that secret information generated in the protocol cannot subsequently be made public [15], this also shows that the generated key bit is composably secure.

Although the protocol generates only a single key bit, it can be simply modified to generate more key bits, still using only two devices, based on correlations introduced in [14]. The modified protocol uses devices with $L > 2$ outcomes on each side and I_N is replaced by the quantity

$$I_{N,L}(P_{XY|AB}) := P(X \oplus_L 1 \neq Y|A = 0, B = 2N-1) + \sum_{\substack{a,b \\ |a-b|=1}} P(X \neq Y|A = a, B = b),$$

where $X \oplus_L 1$ represents addition modulo L. This protocol can be implemented by quantum devices containing maximally entangled L-dimensional quantum states and carrying out measurements with L possible outcomes [14].

The next section contains a precise statement and proof of security for Protocol R.

Security definition

We use here a standard definition of composable security (based on the definitions in [20], previously applied in an analogous way to our treatment in [11, 12, 21]). A composable security definition should ensure that a protocol is not only secure for a single instance, but that it remains secure if used as a sub-protocol in part of an arbitrary extended protocol. In order to show this, one considers an ideal protocol (that is by definition secure) and proves that there is no extended protocol that can correctly guess whether it is interfacing with the ideal or real

[2] A note about that notation used in this paper. We tend to use upper case for random variables, and lower case for particular instances of them. In addition, $P_{X|A=a}$ is the distribution over the random variable X conditioned on the event that random variable A takes value a. This will often be abbreviated to $P_{X|a}$. There is another common notation in which this is written $P(X|A = a)$.

protocol with probability significantly greater than $\frac{1}{2}$. Roughly speaking, the idea is that if this holds, the two protocols behave essentially identically when used as part of any other protocol. Furthermore, if the probability of correctly guessing differs from $\frac{1}{2}$ by at most p, then, for n uses of either the real protocol or idea, the probability of correctly guessing differs from $\frac{1}{2}$ by at most np.

Formally, one considers a *distinguisher*, that tries to guess which protocol (the real or ideal) is being used. For two key distribution protocols, 1 and 2, a distinguisher is an extended protocol that uses the candidate protocol as a sub-protocol, and outputs a single bit, corresponding to a guess of whether the sub-protocol was protocol 1 or 2. The distinguisher can ask the eavesdropper to act in any way[3]. and can use Eve's outputs, those of the honest parties and any information made public in the protocol's implementation to try to distinguish the two. It does not, however, have access to any private data that the honest users use.

Let us denote by Γ the complete set of random variables the distinguisher receives from Alice and Bob during the protocol, as well as the protocol's outputs. If Protocol 1 is followed, these are distributed according to Q^1_Γ, while if Protocol 2 is followed, these are distributed according to Q^2_Γ (for some fixed device behaviour chosen by Eve). Having received these, the distinguisher has access to a system (held by Eve) with input denoted C, and output Z. The probability of correctly guessing whether Alice and Bob are following Protocol 1 or 2 (chosen with probability $\frac{1}{2}$ each) is given by[4].

$$\frac{1}{2}(1+\frac{1}{2}\sum_\gamma \max_c \sum_z |Q^1_\Gamma(\gamma)Q^1_{Z|\gamma c}(z) - Q^2_\Gamma(\gamma)Q^2_{Z|\gamma c}(z)|)\,. \tag{3}$$

The notion of security we use is based on the success probability of the optimal distinguisher (i.e. where the distinguisher asks Eve to behave in such a way as to make distinguishing easiest)[5].

[3] Note that what Eve does can be adapted depending on any information available to the distinguisher.

[4] Note that in the case that Eve keeps only a classical system (so there is no c), this reduces to $\frac{1}{2}(1 + D(P^1_{\Gamma Z}, P^2_{\Gamma Z}))$, where D denotes the *total variation distance* (defined later).

[5] A note on notation: we characterize the behaviour of the devices by the joint conditional probabilities of the outputs if the inputs are chosen independently, and label these using P. For example, in the case of three devices shared between Alice, Bob and Eve, these are denoted $P_{XYZ|ABC}$ and are assumed to satisfy the no-signalling conditions (1). We use expressions involving Q (e.g., $Q_{\Gamma CZ}$) to denote the actual distribution of random variables in the scenario where a protocol is being performed on these systems in conjunction with a distinguisher. There is an important distinction between the two: since a distinguisher can arrange that C is correlated with Γ, Q may no longer obey the no-signalling conditions (1). For example, if Γ includes the output, X, of Alice's device (whose input is A), and the distinguisher chooses $C = X$, the non-signalling condition $Q_{X|AC} = Q_{X|A}$ does not generally hold.

Definition 1. Protocol 1 is said to be ζ-*secure* with respect to Protocol 2 if the probability of correctly guessing whether a candidate protocol is Protocol 1 or 2 (chosen with probability $\frac{1}{2}$ each) by any distinguisher is at most $\frac{1}{2}(1+\zeta)$.

We define an ideal protocol, Protocol ID, to be identical to Protocol R, except that Step 5 is replaced by

5′. If $K = 1$, Alice and Bob take their outputs from a hypothetical device that gives X to Alice, and Y to Bob such that $X = Y$ and X is uniformly distributed and uncorrelated with any other information.

This protocol either aborts (with the same probability as Protocol R), or outputs the same perfectly private bit to both Alice and Bob.

In order to prove security of Protocol R, it is useful to define a modified protocol, to be used as a technical tool in the proof. We consider a protocol that is the same as Protocol R, except with a more powerful eavesdropper who, before the protocol restarts at the end of Step 4, has access to all the data previously produced and can alter Alice's and Bob's devices at this stage. Formally, let Protocol R′ be identical to Protocol R, except that Step 4 is replaced by

4′. If $K = 0$, Alice and Bob publicly announce their outputs X'_i and Y_i. If $X'_i \neq Y_i$, they abort. Otherwise, they return their devices to Eve who can modify them and supply new ones. Alice and Bob both announce receipt of their new devices, before returning to Step 1.

We also define an analogous ideal, Protocol ID′, which is obtained from Protocol ID by replacing Step 4 with Step 4′.

The reason for this adjustment is that Protocol R′ clearly cannot be more secure than Protocol R (the set of allowed actions of Eve in Protocol R′ is strictly larger than that in Protocol R). Hence it is sufficient to prove security of Protocol R′. But the analysis of Protocol R′ is relatively simple, because the optimal distinguisher will ask the eavesdropper to act in an independent and identically distributed (i.i.d.) way on each round, and is essentially characterized by the single constant value of I_N used on each round.

We will show that Protocol R′ is ζ-secure with respect to Protocol ID′, where the parameter ζ can be made arbitrarily small by appropriate choices of α and N. Since both protocols have identical probabilities of aborting, an abort event cannot help the distinguisher. Furthermore, in any strategy with a significant probability of not aborting, the protocols remain virtually indistinguishable. This shows that Protocol R′ is composably secure in the appropriate sense. As mentioned before, it follows that Protocol R is also ζ-secure with respect to Protocol ID and hence also composably secure.

Security proof

The proof bounds the probability of distinguishing Protocols R′ and ID′. First, note that there is an optimal distinguishing strategy in which Eve's actions are i.i.d. since, if it does not abort, when the protocol returns to Step 1, the maximum probability of distinguishing the protocols is identical to that before the protocol began.

We use the following lemma, that uses I_N to bound the distance between probability distributions, measured using the total variation distance, $D(P_X, Q_X) := \frac{1}{2}\sum_x |P_X(x) - Q_X(x)|$. The proof of this lemma can be found in [22, Supplementary Information] (and is based on similar results in [5, 14, 23]):

Lemma 2. *[22] For any non-signalling device behaviour, $P_{XYZ|ABC}$, in which X and Y are binary, we have*

$$D(P_{Z|abcx}, P_{Z|c}) \leq I_N(P_{XY|AB}) \tag{4}$$

for all a, b, c and x, and

$$D(P_{X|abc}, \frac{1}{2}) \leq \frac{1}{2} I_N(P_{XY|AB}) \tag{5}$$

for all a, b and c.

(Note: We use $D(P_{X|abc}, \frac{1}{2})$ to denote the distance between $P_{X|abc}$ and the distribution where $X = 0$ and $X = 1$ both occur with probability $\frac{1}{2}$.)

Note that these relations imply

$$D(P_{Z|abc,X'=x}, P_{Z|c}) \leq I_N(P_{XY|AB}) \tag{6}$$

and

$$D(P_{X'|ab}, \frac{1}{2}) \leq \frac{1}{2} I_N(P_{XY|AB})\,. \tag{7}$$

Note also that, from the definition of I_N, averaging over the measurements in $\mathcal{G}_N$ (picked uniformly), we have

$$\begin{aligned} P(X' \neq Y) &:= \sum_{\substack{abx \\ (a,b)\in\mathcal{G}_N}} \frac{P_{X'Y|ab}(x, 1-x)}{2N} \\ &= \frac{I_N(P_{XY|AB})}{2N}\,. \end{aligned} \tag{8}$$

We also need the following generalization of (6):

Lemma 3. *For any non-signalling device behaviour, $P_{XYZ|ABC}$, in which X and Y are binary, and $I_N := I_N(P_{XY|AB}) < 1$ we have that, for $(a, b) \in \mathcal{G}_N$,*

$$D(P_{Z|abc,X'=x,Y=x}, P_{Z|abc,X'=x}) \leq \frac{2I_N}{1 - I_N}\,. \tag{9}$$

Proof. We have

$$\begin{aligned} P_{Z|abc,X'=x,Y=x}(z) - P_{Z|abc,X'=x}(z) &= P_{Z|abc,X'=x,Y=x}(z) - \sum_y P_{YZ|abc,X'=x}(y,z) \\ &= P_{Z|abc,X'=x,Y=x}(z) - P_{Y|abc,X'=x}(x) P_{Z|abc,X'=x,Y=x}(z) - P_{Y|abc,X'=x}(1-x) P_{Z|abc,X'=x,Y=1-x}(z) \\ &= (1 - P_{Y|abc,X'=x}(x))(P_{Z|abc,X'=x,Y=x}(z) - P_{Z|abc,X'=x,Y=1-x}(z)) \end{aligned}$$

and hence

$$\begin{aligned} D(P_{Z|abc,X'=x,Y=x}, P_{Z|abc,X'=x}) &= (1 - P_{Y|abc,X'=x}(x)) D(P_{Z|abc,X'=x,Y=x}(z), P_{Z|abc,X'=x,Y=1-x}(z)) \\ &\leq (1 - P_{Y|abc,X'=x}(x))\,. \end{aligned}$$

Then note that averaging over the measurements in $\mathcal{G}_N$, using (8) we have

$$\begin{aligned} 1 - \frac{I_N}{2N} &= \sum_{\substack{a'b'x \\ (a',b')\in\mathcal{G}_N}} \frac{1}{2N} P_{X'Y|a'b'}(x,x) \\ &\leq \frac{1}{2N}\left(\sum_x P_{X'Y|ab}(x,x) + 2N - 1\right), \end{aligned}$$

from which it follows that

$$\sum_x P_{X'Y|ab}(x,x) \geq 1 - I_N\,,$$

and hence

$$\begin{aligned} P_{Y|abc,X'=x}(x) &= \frac{P_{X'Y|abc}(x,x)}{P_{X'|abc}(x)} \\ &\geq \frac{1}{P_{X'|ab}(x)}\left(1 - I_N - P_{X'Y|abc}(1-x, 1-x)\right) \\ &\geq \frac{1}{P_{X'|ab}(x)}\left(1 - I_N - (1 - P_{X'|ab}(x))\right) \\ &\geq 1 - \frac{2I_N}{1 - I_N}\,, \end{aligned}$$

where we used (7) in the last line. Note that the last step does not hold unless $I_N < 1$. The claimed relation then

follows. □

Combining (9) and (6) (using the triangle inequality for D), we have for $I_N < 1$

$$D(P_{Z|abc,X'=x,Y=x}, P_{Z|c}) \leq (1 + \frac{2}{1 - I_N})I_N \,. \qquad (10)$$

To successfully distinguish the protocols it is necessary that they do not abort before the final round. We use $\perp$ to represent the event that the protocol aborts, and $\bar{\perp}$ to represent the event that it does not.

Lemma 4. *For $0 \leq I_N^* \leq 2N$, if Protocol R' is followed, and Eve supplies i.i.d. states corresponding to non-signalling device behaviours with $I_N(P_{X_iY_i|A_iB_i}) = I_N^*$ for all i, then*

$$Q(\bar{\perp}) = \left(1 + \frac{(1-\alpha)I_N^*}{2N\alpha}\right)^{-1} .$$

Proof. We have

$$Q(f = j) = \left((1-\alpha)(1 - \frac{I_N^*}{2N})\right)^{j-1} \alpha \,,$$

and hence

$$Q(\bar{\perp}) = \sum_{j=1}^{\infty} Q(f = j) = \left(1 + \frac{(1-\alpha)I_N^*}{2N\alpha}\right)^{-1} ,$$

as required. □

Our main result is then as follows

Theorem 5. *Take $\alpha = N^{-\frac{3}{2}}$. Then Protocol R' is ζ-secure with respect to ID' for $\zeta = \frac{23}{2}N^{-1/2}$. Furthermore, in a noise-free implementation with honest devices, Protocol R' does not abort with probability greater than $(1 + \pi^2 N^{-1/2}/16)^{-1}$.*

Proof. As mentioned above, Protocols R' and ID' can be optimally distinguished when the eavesdropper supplies i.i.d. states, and so her device behaviour can be characterized by a single value, I_N^*, the value of $I_N(P_{X_iY_i|A_iB_i})$ on each round i. The two protocols are identical up to Step 5, and so can be distinguished only if the protocol does not abort. In the case of no abort, the distinguisher sees A_f, B_f, X'_f and Y_f, and then has access to a system with input C and output Z. (The distinguisher also has data from previous rounds, but these are identically distributed for Protocols R' and ID', and so can be ignored.) Noting that the device behaviour of the ideal obeys

$$P^{\text{ID}'}_{X'YZ|abc}(x, y, z) := \frac{1}{2}\delta_{x,y} P^{\text{R}'}_{Z|c}(z) \,,$$

we can relate the terms in (3) to the device behaviours of the real and ideal as follows:

$$\begin{aligned}
Q^{\text{R}'}_{A_fB_fX'_fY_f} &= \frac{1}{2N} P^{\text{R}'}_{X'_fY_f|A_fB_f} \\
Q^{\text{R}'}_{Z|A_fB_fCX'_fY_f} &= P^{\text{R}'}_{Z|A_fB_fCX'_fY_f} \\
Q^{\text{ID}'}_{A_fB_fX'_fY_f} &= \frac{1}{4N}\delta_{x,y} \\
Q^{\text{ID}'}_{Z|A_fB_fCX'_fY_f} &= P^{\text{R}'}_{Z|C} \,.
\end{aligned}$$

For convenience, we drop the subscript f in the following.

We will consider two separate cases. The first is $I_N^* \geq 1/2$. In this case, we can upper bound the probability of correctly distinguishing the protocols by assuming that they can be perfectly distinguished in the case that the protocol does not abort. Using Lemma (4), it follows that in this case the probability of correctly guessing which protocol is being used can be upper bounded by

$$\frac{1}{2}(1 + Q(\bar{\perp})) \leq \frac{1}{2}(1 + 4N^{-\frac{1}{2}}) \,,$$

where we have substituted the value of α and used

$$(1 + 4N^{-\frac{1}{2}} - N^{-\frac{3}{2}})^{-1} \leq 1 \qquad (11)$$

for $N \geq 2$ to simplify the bound.

Turning now to the case $I_N^* \leq 1/2$, the probability of correctly guessing which protocol is being followed is $\frac{1}{2}(1 + Q(\bar{\perp})\Delta)$, where

$$\begin{aligned}
\Delta &:= \frac{1}{2} \sum_{\substack{a,b,x,y \\ (a,b)\in\mathcal{G}_N}} \max_c \sum_z |Q^{\text{R}'}_{ABX'Y} Q^{\text{R}'}_{Z|ABCX'Y} - Q^{\text{ID}'}_{ABX'Y} Q^{\text{ID}'}_{Z|ABCX'Y}| \\
&= \frac{1}{4N} \sum_{\substack{a,b,x,y \\ (a,b)\in\mathcal{G}_N}} \max_c \sum_z |P^{\text{R}'}_{X'Y|ab}(x,y) P^{\text{R}'}_{Z|abcxy}(z) - \frac{1}{2}\delta_{x,y} P^{\text{R}'}_{Z|c}(z)| \\
&= \frac{1}{4N} \sum_{\substack{a,b,x,y \\ (a,b)\in\mathcal{G}_N, x=y}} \max_c \sum_z |P^{\text{R}'}_{X'Y|ab}(x,y) P^{\text{R}'}_{Z|abcxy}(z) - \frac{1}{2} P^{\text{R}'}_{Z|c}(z)| + \frac{1}{4N} \sum_{\substack{a,b,x,y \\ (a,b)\in\mathcal{G}_N, x\neq y}} \max_c \sum_z P^{\text{R}'}_{X'Y|ab}(x,y) P^{\text{R}'}_{Z|abcxy}(z) \,.
\end{aligned}$$

The second term is equal to $\frac{1}{4N}\sum_{\substack{a,b,x,y\\(a,b)\in\mathcal{G}_N,x\neq y}} P^{\mathrm{R}'}_{X'Y|ab}(x,y) = \frac{1}{2}P^{\mathrm{R}'}(X'\neq Y)$, and the first term is equal to

$$\frac{1}{4N}\sum_{\substack{a,b,x,y\\(a,b)\in\mathcal{G}_N,x=y}} \max_c \sum_z |P^{\mathrm{R}'}_{X'Y|ab}(x,y)P^{\mathrm{R}'}_{Z|abcxy}(z) - \frac{1}{2}P^{\mathrm{R}'}_{Z|c}(z)|\,.$$

Then note that

$$\begin{aligned}\sum_z |P^{\mathrm{R}'}_{X'Y|ab}(x,x)P^{\mathrm{R}'}_{Z|abcxy}(z) - \frac{1}{2}P^{\mathrm{R}'}_{Z|c}(z)| \;&\leq\; \sum_z |P^{\mathrm{R}'}_{X'Y|ab}(x,x)P^{\mathrm{R}'}_{Z|abcxy}(z) - P^{\mathrm{R}'}_{X'Y|ab}(x,x)P^{\mathrm{R}'}_{Z|c}(z)|\\ &\quad + \sum_z |P^{\mathrm{R}'}_{X'Y|ab}(x,x)P^{\mathrm{R}'}_{Z|c}(z) - \frac{1}{2}P^{\mathrm{R}'}_{Z|c}(z)|\\ &=\; \sum_z P^{\mathrm{R}'}_{X'Y|ab}(x,x)|P^{\mathrm{R}'}_{Z|abcxx}(z) - P^{\mathrm{R}'}_{Z|c}(z)| + \sum_z P^{\mathrm{R}'}_{Z|c}(z)|P^{\mathrm{R}'}_{X'Y|ab}(x,x) - \frac{1}{2}|\\ &\leq\; 2P^{\mathrm{R}'}_{X'Y|ab}(x,x)(\frac{2I^*_N}{1-I^*_N} + I^*_N) + |P^{\mathrm{R}'}_{X'Y|ab}(x,x) - \frac{1}{2}|\,.\end{aligned}$$

where we have used (10). In addition,

$$\begin{aligned}|P^{\mathrm{R}'}_{X'Y|ab}(x,x) - \frac{1}{2}| \;&\leq\; |P^{\mathrm{R}'}_{X'Y|ab}(x,x) - P^{\mathrm{R}'}_{X'|ab}(x)| + |P^{\mathrm{R}'}_{X'|ab}(x) - \frac{1}{2}|\\ &=\; P^{\mathrm{R}'}_{X'|ab}(x) - P^{\mathrm{R}'}_{X'Y|ab}(x,x) + |P^{\mathrm{R}'}_{X'|ab}(x) - \frac{1}{2}|\,.\end{aligned}$$

Bringing everything together, we have

$$\begin{aligned}\Delta &\leq \frac{1}{4N}\sum_{\substack{a,b,x\\(a,b)\in\mathcal{G}_N}} \left(P^{\mathrm{R}'}_{X'Y|ab}(x,x)(\frac{4}{1-I^*_N}+2)I^*_N + P^{\mathrm{R}'}_{X'|ab}(x) - P^{\mathrm{R}'}_{X'Y|ab}(x,x) + |P^{\mathrm{R}'}_{X'|ab}(x) - \frac{1}{2}|\right) + \frac{1}{2}P^{\mathrm{R}'}(X'\neq Y)\\ &\leq (\frac{2}{1-I^*_N}+1)I^*_N + \frac{1}{2}(1-P^{\mathrm{R}'}(X'=Y)) + \frac{I^*_N}{2} + \frac{1}{2}P^{\mathrm{R}'}(X'\neq Y) = \left(\frac{2}{1-I^*_N}+\frac{3}{2}+\frac{1}{2N}\right)I^*_N \leq \frac{23}{4}I^*_N\,,\end{aligned}$$

where we used (7), (8), and the last bound relies on $I^*_N \leq 1/2$ and $N \geq 2$. The distinguisher's probability of correctly guessing is thus

$$\frac{1}{2}(1+Q(\bar{\perp})\Delta) \leq \frac{1}{2}\left(1+\left(1+\frac{(1-\alpha)I^*_N}{2N\alpha}\right)^{-1}\frac{23}{4}I^*_N\right).$$

Maximizing over $0 \leq I^*_N \leq 1/2$ gives a maximum of $\frac{1}{2}(1+\frac{23N\alpha}{2(4N\alpha+1-\alpha)})$ at $I^*_N = 1/2$. Substituting $\alpha = N^{-\frac{3}{2}}$ and using (11), we can upper bound this by $\frac{1}{2}(1+\frac{23}{2}N^{-\frac{1}{2}})$. Since we have already established a tighter bound for $I^*_N \geq \frac{1}{2}$, this completes the first part of the claim.

The probability of an abort in the case that Eve supplies honest devices (and there is no noise) can be calculated as in Lemma 4, except that in this situation, each round has $I_N = I^{\mathrm{QM}}_N < \pi^2/8N$ (cf. (2)). The probability that the protocol does not abort is then

$$\left(1+\frac{(1-\alpha)I^{\mathrm{QM}}_N}{2N\alpha}\right)^{-1} > \left(1+\frac{\pi^2}{16N^2\alpha}\right)^{-1}, \qquad (12)$$

from which the claim is recovered by substituting the value of α. □

For sufficiently large N, we can hence make ζ as close to 0 as we like, at the same time as making the probability of an abort in the absence of Eve close to 0.

Finally, since, by construction, it is harder to distinguish Protocol R from Protocol ID than it is to distinguish Protocol R′ from ID′, Protocol R is also ζ-secure with respect to Protocol ID for the same ζ, and an analogous statement can be made about Protocol R^+. Clearly, too, when Eve is honest and noise is absent, Protocols R, R′ and R^+ all have the same abort probability, in each case bounded by (12).

Attacks on modified protocols by a post-quantum eavesdropper

Protocol R relies on a probabilistic strategy in which Alice and Bob sequentially either (with high probability) test a purported entangled state generated by their devices or (with low probability) generate a key bit from the state and immediately end the protocol. We consider below two seemingly natural modifications of Protocol R and highlight some interesting attacks available

to an eavesdropper in such cases. The first of our modified protocols can be broken by a quantum eavesdropper and that the second can be broken by an eavesdropper restricted only by signalling constraints.

Protocol S

This protocol is specified by positive integers M and N.

1. On the i^{th} round, Alice picks a pair of values (A_i, B_i) at random from the set $\mathcal{G}_N$, and announces B_i to Bob.

2. Alice inputs A_i into her device, and Bob B_i into his, and they record their outcomes, the bits X_i and Y_i respectively. (Alice ensures that her device doesn't learn B_i.) If $(A_i, B_i) = (0, 2N-1)$, Alice sets $X_i' = 1 - X_i$, otherwise she sets $X_i' = X_i$. The protocol returns to Step 1 unless $i = M$.

3. Alice randomly chooses an integer $1 \leq f \leq M$ and announces it to Bob.

4. Alice and Bob publicly announce X_i' and Y_i for all $i \neq f$. If any of their announced values are unequal, they abort.

5. The bits X_f' and Y_f are taken to be the final shared bit.

This protocol is similar in spirit to the original BHK protocol [5], and vulnerable to the same kind of attack in the scenario where Alice and Bob have only one device each.

In this case, if Eve equips her devices with memory, she has a simple attack. She programs her devices to behave honestly until the final (M^{th}) round. On this round, Alice's device outputs the XOR of the previous outputs, i.e. $\bigoplus_{i=1}^{M} X_i$, and Bob's device outputs a random bit. This attack leads to a probability of abort close to $\frac{1}{2}$, and otherwise enables Eve to perfectly guess the final output bit. Crucially, the success probability of this strategy cannot be made small by adjusting M and N.

Define Protocol T by altering Step 4 of Protocol S to circumvent this attack:

$\bar{4}$. For all $i \neq f$, Alice chooses $L_i = 0$ with probability β and $L_i = 1$ with probability $1-\beta$. She announces this list to Bob. For all the rounds in which $L_i = 1$, Alice and Bob publicly announce their outcomes. If any of their announced values are unequal, they abort.

With this modification, making the final output the XOR of the previous ones does not give Eve significant information, since Eve no longer learns all but one of the outputs, $\{X_i\}$. However, there is another attack that a post-quantum non-signalling eavesdropper can use in this case, which allows her to learn the final bit, again with a probability of success that cannot be made small for any choice of M and N. This attack exploits some subtle properties of non-local correlations and cannot be performed by a quantum-limited eavesdropper.

The attack is based on a result in [24] and involves non-local boxes [25, 26]. These are bipartite systems where each party has two choices of input and receives one of two outputs. If we denote the inputs $x \in \{0,1\}$ and $z \in \{0,1\}$ and the respective outputs $\alpha \in \{0,1\}$ and $\gamma \in \{0,1\}$, then the non-local box is a non-signalling device which outputs according to $x.z = \alpha \oplus \gamma$.

The attack is as follows. Eve constructs Alice's device such that it contains both a set of maximally entangled quantum states shared with Bob, and a set of non-local boxes shared with Eve (the same number of each). For the first $M - \frac{1}{\beta}$ rounds of the protocol, Alice's device generates its output by making quantum measurements as in an honest implementation of the protocol. However, as well as supplying the measurement outcome to the output port of the device (so that Alice sees it), the outcome is also used as input to one of the non-local boxes, generating an output (call it α_i). (Bob's device behaves honestly in the first $M - \frac{1}{\beta}$ rounds, and outputs predetermined random bits in the remaining ones.)

In the last $\frac{1}{\beta}$ rounds, Alice's device instead always outputs the XOR of all the previous non-local box outputs, i.e. $\bigoplus_i \alpha_i$. (Although this may look suspicious, it does not violate the stated security tests. In any case it could easily be masked by shared randomness between Alice's device and Eve.) With reasonable probability, Eve will learn this bit (on each round of the protocol, the chances that the output of that round is communicated between Alice and Bob is β, so, of the last $\frac{1}{\beta}$, on average 1 will be communicated). For each bit of the last $\frac{1}{\beta}$ that is communicated there is a probability 1/2 of being detected by Alice and Bob, so this strategy implies a significant probability that Eve will be detected. However, the probability that this attack works without detection is independent of N and M and at least $\frac{1}{2e}$.

If Eve learns $\bigoplus_i \alpha_i$, she can determine the key bit, x_f. To see this, notice the non-local box condition is $x_i.z_i = \alpha_i \oplus \gamma_i$, where z_i are the inputs and γ_i the outputs of Eve's half of the non-local box. Eve should input 0 to all of her halves of the non-local boxes, except the f^{th} one in which she inputs 1. We have

$$x_f = \bigoplus_i (x_i.z_i) = \bigoplus_i (\alpha_i \oplus \gamma_i) = \bigoplus_i \alpha_i + \bigoplus_i \gamma_i.$$

Therefore, provided she has obtained the bit $\bigoplus_i \alpha_i$, Eve can determine the final bit output by the protocol, x_f.

Attacking more noise-tolerant protocols

In this section, we consider some extensions of the type of attack considered in the previous section to two-device protocols that (if secure) would be more efficient and tolerate more noise.

In all device-independent key distribution protocols, one needs, in essence, to establish the presence of non-local correlations. In order to do so, the detection loophole must be closed. In other words, a malicious device should not be able to exploit detector failures (cases where no outcome is observed) to give the false illusion of non-locality in the non-failure cases.

Protocols based on chained Bell correlations with large N, are not well-suited to this, since as N increases, it becomes increasingly difficult to close the detection loophole (the correlations can be classically explained if the probability of detector failure is $\frac{1}{N}$). This drawback is not limited to the two-device case, and alternative protocols tolerating modest levels of noise have been introduced in the case where more devices are permitted [12, 13]. We now consider the extension of these protocols to the two-device case. We do not give a proof that all such protocols are insecure, but give an example that highlights interesting security issues that can arise in the presence of non-signalling eavesdroppers.

We also mention some other work related to this question. In [16], the two-device case was considered for protocols based on CHSH correlations. There it was shown that privacy amplification via hashing is not possible against an adversary limited only by the impossibility of signalling between the parties. However, in [16], signalling was permitted within the devices (so that outputs could depend on later inputs[6]). For protocols in which each party waits for an output before giving their next input, the most natural signalling constraints are ones that allow later outputs to depend on all previous inputs, but do not allow outputs to depend on future inputs (we call these *time-ordered non-signalling conditions*). A situation that is close to this case (but with subtle and potentially important differences) has been recently studied in [17]. There protocols based on CHSH correlations were again considered, and it was shown that privacy amplification via hashing is not possible for adversaries limited by almost time-ordered non-signalling conditions.

Consider now a key distribution protocol with the following structure[7]:

1. Alice and Bob each make a random input A_i and B_i to their devices, ensuring they receive their outputs (X_i and Y_i respectively) before making the next input (so that time-ordered non-signalling conditions must be obeyed). They repeat this M times.

2. Either Alice, or Bob (or both) publicly announces their measurement choices, and one party checks that they had a sufficient number of the relevant input combinations, and otherwise aborts. Certain rounds may be discarded according to some public protocol.

3. For each of the remaining bits, Alice independently announces it to Bob with probability μ (which is such that $M\mu$ is large). Bob uses this to compute some test function. If this has the wrong output, Bob aborts. (For example, Bob might compute the CHSH value of the announced data, and abort if it is below 2.5.) (This step is often called *parameter estimation*.)

4. Alice and Bob perform error correction using public communication via any protocol in which the function Alice applies to her string becomes known to Eve[8].

5. Alice and Bob publicly perform privacy amplification. The function Alice applies to her string becomes known to Eve[8].

$P_{SZ\|XC}$	C	0		1		2	
	Z	0	1	0	1	0	1
X	S						
$00\ldots00$	0	1/2	0	0	1/2	$\ldots$	
	1	0	1/2	1/2	0		
$00\ldots01$	0	0	1/2	1/2	0		
	1	1/2	0	0	1/2		
	$\vdots$			$\vdots$			

TABLE I: **Behaviour of the "joint function box".** Each 2×2 block takes one of the two forms shown, depending on whether $F_C(X_1 \ldots X_M) = 0$ or $F_C(X_1 \ldots X_M) = 1$. In this notation, the non-signalling conditions are that the sum of the elements in each row of each 2×2 block are equal to those of the blocks to the left and right, and likewise, the sum of the elements in each column are equal to those above and below. In the above case, all of these values are 1/2.

The key to Eve's attack is Step 3. She is going to attack so as to try to gain one bit of the final output string. Eve will also use a "joint function box", which has the following bipartite behaviour. Alice inputs a string $X_1 \ldots X_M$ and obtains a single bit S, Eve inputs C which corresponds to a choice of one-bit function (see later), and obtains a single bit Z. The behaviour is such that $Z = S \oplus F_C(X_1 \ldots X_M)$. It is easy to see that this can be non-signalling if S is a uniform random bit.

It follows that Eve can learn any Boolean function of $X_1, \ldots, X_M$ if she receives just one bit, S. The value of C depends on the function Eve wants to learn, the value of S she hears and the information reconciliation and privacy amplification functions she overhears. There is a choice of C for each combination of these values. Thus,

[6] Although, as currently described, this is unphysical, it is natural to consider this for protocols in which each party makes all their inputs at the start, and then receives all of their outputs together.

[7] Although this structure is not fully general, most protocols to date are of this type.

[8] Typically because it is communicated over the public channel.

for protocols of the above form (importantly, where Eve learns the entire function Alice and Bob use for post-processing), she needs to receive only one bit from either of her devices to learn one bit about the final output key (after privacy amplification).

In order to try to learn this bit, Eve can exploit the parameter estimation step. She programs Alice's device to behave honestly for the first $M - 1/\mu$ rounds (note that we have not specified which correlations are used; this attack does not depend on these (up to a constant factor in the abort probability) and even works if the honest states are perfect non-local boxes). Her device then inputs the bits generated in these rounds into the "joint function box", producing output S. This bit is then given as the outputs X_i for $M - 1/\mu \leq i \leq M$ (this could be masked by XORing with some pre-shared randomness between Alice's device and Eve). Provided at least one of the last $1/\mu$ bits is revealed in the parameter estimation without causing abort (this occurs with finite probability that cannot be made arbitrarily small by judiciously choosing M and μ), Eve can discover any desired bit of the final output string.

There are a couple of important points to note about the above attack. Firstly, we assumed a specific protocol structure. In particular, altering the way parameter estimation is done could potentially improve security (some altered protocols are discussed in [15]). Secondly, the attack relies on a specialized non-local strategy that cannot be implemented by an eavesdropper limited by quantum theory. Proving security of a protocol of this type (in particular, with two devices) that is secure against a quantum-restricted Eve remains an open question.

Conclusions

We have presented a protocol for distribution of a one-bit key, and have proven it secure in a universally composable way against an arbitrarily powerful adversary who can create all the supposedly quantum devices, provided that the devices are not reused in any future protocol. The protocol only requires two devices, whereas the secure protocols previously considered required many independent devices. This represents a theoretical advance, and also potentially represents another step towards practical unconditionally secure device-independent key distribution protocols.

That said, several significant and intriguing theoretical and practical issues remain. First, the simplest version of our protocol only outputs a single bit, requiring a large number of entangled qubit pairs in order to do so. The protocol can be generalised to produce an arbitrary length key string, but again, highly inefficiently. It would be very interesting to know whether significantly more efficient two-device secure protocols can be found, and to obtain bounds on what is achievable.

Secondly, for maximum flexibility and more efficient use of resources, one would like to be able to repeat the protocol to generate further secure key bits. However, if devices are reused, this renders the protocol vulnerable to the same device-memory-based attacks [15] that apply to BHK and other device-independent protocols. While it is clear that device-reusing protocols cannot be universally composable, the general scope of such attacks and the possibilities of countering them either by refined protocols (see [15] for ideas in this direction, some of which have been later developed in [27]) or by evidently reliable technological assumptions have not yet been fully explored. It would be very interesting to resolve these questions in the present context.

Thirdly, tolerance to noise is a significant practical issue for our protocol. As given, it aborts if there is one set of measurements that give unequal outcomes. The protocol parameters are tuned such that this is very unlikely if the devices operate perfectly. However, with more realistic, noisy devices, using present technology, the protocol will have a very high abort rate. Although the protocol could be adapted to tolerate small amounts of noise, it is far from being practical in this respect.

Fourthly, our scheme requires an authenticated (although public) classical channel, and, a common way to implement this in an information-theoretically secure way using an insecure classical channel, is by using a pre-shared key. This reinforces the points already made: it would be desirable to have more efficient two-device protocols that allow for some consumption of key for classical authentication and nonetheless provide quantum key expansion at practically useful rates in realistically noisy environments.

In summary, while we have presented a protocol showing that device-independent quantum key distribution is in principle possible using two devices, a number of theoretically interesting and practically important questions remain open.

Remark: In some concurrent work an alternative technique for proving security of device-independent QKD with two devices has been suggested [28]. Furthermore, since the first version of our paper, an additional article has appeared [29] reporting an efficient and noise-tolerant scheme. We note that these works differ from ours in that they consider quantum-limited eavesdroppers and do not apply to the case of eavesdroppers limited only by signalling constraints.

Acknowledgments

We thank Graeme Mitchison for many useful conversations. JB was supported by the EPSRC, and the CHIST-ERA DIQIP project. RC acknowledges support from the Swiss National Science Foundation (grants PP00P2-128455 and 20CH21-138799) and the National Centre of Competence in Research 'Quantum Science and Technology'. AK was partially supported by a Leverhulme Research Fellowship, a grant from the John Templeton

Foundation, and the EU Quantum Computer Science project (contract 255961). This research is supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation.

[1] Bennett, C. H. & Brassard, G. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, 175–179. IEEE (New York, 1984).

[2] Ekert, A. K. Quantum cryptography based on Bell's theorem. *Physical Review Letters* **67**, 661–663 (1991).

[3] Gerhardt, I. *et al.* Full-field implementation of a perfect eavesdropper on a quantum cryptography system. *Nature Communications* **2**, 349 (2011).

[4] Mayers, D. & Yao, A. Quantum cryptography with imperfect apparatus. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS-98)*, 503–509 (IEEE Computer Society, Los Alamitos, CA, USA, 1998).

[5] Barrett, J., Hardy, L. & Kent, A. No signalling and quantum key distribution. *Physical Review Letters* **95**, 010503 (2005).

[6] Acin, A., Gisin, N. & Masanes, L. From Bell's theorem to secure quantum key distribution. *Physical Review Letters* **97**, 120405 (2006).

[7] Scarani, V. *et al.* Secrecy extraction from no-signaling correlations. *Physical Review A* **74**, 042339 (2006).

[8] Acin, A. *et al.* Device-independent security of quantum cryptography against collective attacks. *Physical Review Letters* **98**, 230501 (2007).

[9] Masanes, L., Renner, R., Christandl, M., Winter, A. & Barrett, J. Unconditional security of key distribution from causality constraints. e-print `quant-ph/0606049v4` (2009).

[10] Masanes, L. Universally composable privacy amplification from causality constraints. *Physical Review Letters* **102**, 140501 (2009).

[11] Hänggi, E., Renner, R. & Wolf, S. Quantum cryptography based solely on Bell's theorem. In Gilbert, H. (ed.) *Proceedings of the 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques (Eurocrypt'10)*, 216–234 (Springer, 2010). Also available `arXiv:0911.4171`.

[12] Hänggi, E. & Renner, R. Device-independent quantum key distribution with commuting measurements. e-print `arXiv:1009.1833` (2010).

[13] Masanes, L., Pironio, S. & Acín, A. Secure device-independent quantum key distribution with causally independent measurement devices. *Nature Communications* **2**, 238 (2011).

[14] Barrett, J., Kent, A. & Pironio, S. Maximally non-local and monogamous quantum correlations. *Physical Review Letters* **97**, 170409 (2006).

[15] Barrett, J., Colbeck, R. & Kent, A. Prisoners of their own device: Trojan attacks on device-independent quantum cryptography. e-print `arXiv:1201.4407` (2012).

[16] Hänggi, E., Renner, R. & Wolf, S. The impossibility of non-signalling privacy amplification. e-print `arXiv:0906.4760` (2009).

[17] Friedman, R. A., Hänggi, E. & Ta-Shma, A. Towards the impossibility of non-signalling privacy amplification from time-like ordering constraints. e-print `arXiv:1205.3736` (2012).

[18] Pearle, P. M. Hidden-variable example based upon data rejection. *Physical Review D* **2**, 1418–1425 (1970).

[19] Braunstein, S. L. & Caves, C. M. Wringing out better Bell inequalities. *Annals of Physics* **202**, 22–56 (1990).

[20] Canetti, R. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS-01)*, 136–145 (2001).

[21] Hänggi, E. *Device-Independent Quantum Key Distribution*. Ph.D. thesis, Swiss Federal Institute of Technology, Zurich (2010). Also available as `arXiv:1012.3878`.

[22] Colbeck, R. & Renner, R. No extension of quantum theory can have improved predictive power. *Nature Communications* **2**, 411 (2011).

[23] Colbeck, R. & Renner, R. Hidden variable models for quantum theory cannot have any local part. *Physical Review Letters* **101**, 050403 (2008).

[24] van Dam, W. Implausible consequences of superstrong nonlocality. e-print `quant-ph/0501159` (2005).

[25] Tsirelson, B. Some results and problems on quantum Bell-type inequalities. *Hadronic Journal Supplement* **8**, 329–345 (1993).

[26] Popescu, S. & Rohrlich, D. Quantum nonlocality as an axiom. *Foundations of Physics* **24**, 379–385 (1994).

[27] McKague, M. & Sheridan, L. Reusing devices with memory in device independent quantum key distribution. e-print `arXiv:1209.4696` (2012).

[28] Reichardt, B. W., Unger, F. & Vazirani, U. Classical command of quantum systems via rigidity of CHSH games. e-print `arXiv:1209.0449` (2012).

[29] Vazirani, U. & Vidick, T. Fully device independent quantum key distribution. e-print `arXiv:1210.1810` (2012).